



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
07.04.2010 Bulletin 2010/14

(51) Int Cl.:
G06N 5/02 (2006.01)

(21) Application number: **08017502.9**

(22) Date of filing: **06.10.2008**

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MT NL NO PL PT RO SE SI SK TR
Designated Extension States:
AL BA MK RS

• **Voegelé, Joachim**
64560 Riedstadt (DE)

(74) Representative: **Heselberger, Johannes et al**
Patent- und Rechtsanwälte
Bardehle - Pagenberg - Dost
Altenburg - Geissler
Galileiplatz 1
81679 München (DE)

(71) Applicant: **Software AG**
64297 Darmstadt (DE)


(72) Inventors:
• **Rislow, Gerald, Dr.**
64347 Griesheim (DE)

(54) **Method and monitoring system for the rule-based monitoring of a service-oriented architecture**

(57) The present invention concerns a method for the rule-based monitoring of a component (C1) added to a service-oriented architecture (SOA), the SOA comprising a plurality of components (C2, ..., Cn) being monitored by a plurality of monitoring components (M2, ..., Mn), the method comprising the following steps:
a. selecting a component (C2) from the plurality of com-

ponents (C2, ..., Cn) of the SOA which is similar to the added component (C1) based on a similarity-metric;
b. selecting at least one rule (R2) from the monitoring component (M2) of the selected component (C2), wherein the at least one rule (R2) has already been validated during the operation of the SOA; and
c. adding the at least one rule (R2) to a new monitoring component (M1) of the added component (C1).

Fig. 2



Description

1. Technical Field

[0001] The present invention relates to a method and a monitoring system for the rule-based monitoring of a service-oriented architecture.

2. The Prior Art

[0002] In a service-oriented architecture (SOA), resources are made available to participants in the network as independent services that the participants can access in a standardized way. Whereas most definitions of a SOA use SOAP requests conveyed via HTTP/HTTPS over TCP/IP, a SOA may use any service-based technology. The services interoperate based on formal definitions which are independent from the underlying platform and programming language. The interface definition encapsulates the specific implementations. A SOA is independent of a specific development technology (such as Java and .NET). The software components become very reusable because the interface is standards-compliant and is independent from the underlying implementation of the (web) service logic. For example, a C# (C Sharp) service could be used by a Java application and vice versa.

[0003] Monitoring a SOA is an important and difficult technical task in order to assure that the SOA operates as desired. This applies not only to the development stage of the SOA, when the source codes of the various SOA components repeatedly need to be debugged but also to the deployment and testing phases, for example when exceptional runtime situations must be controlled or the performance of the SOA is to be evaluated for maintenance or other purposes.

[0004] Traditionally, monitoring a software system comprises generating one or more log files according to a set of fixed and predetermined rules of the monitoring system, which allow for understanding the sequence of events as they occurred in the software system. This concept is problematic for application in a SOA, since it only provides a limited amount of flexibility which is in contrast to the flexible nature of a SOA. Furthermore, log files are not applicable to discover complex error situations in advance. On the contrary, log files only allow discovering the reason for a problem after the problem occurred. Consequently, they do not enable taking any counter measures in due time to prevent the problem.

[0005] Since a SOA landscape usually consists of many components that can interact with each other in a complex way and that may dynamically change over time, keeping track of all configurations and the changes thereof is a non-trivial task. It is therefore known in the prior art to provide a centralized registry/repository for a SOA, such as the products "Infravio" or "CentraSite" of applicant, which can simplify and facilitate this task. However, documenting the SOA landscape does not ensure

a smooth and error-free operation. In the EP 1 863 258 applicant already described a system and method for managing web services and how runtime information of the SOA can be gathered and stored in this registry. Whereas the disclosed method and system already substantially facilitates the management of the complex SOA landscape, it is not capable to automatically analyze complex situation.

[0006] Further monitoring systems known in the prior art are bound to application servers and mostly supervise only web services. An example can be found in the US 2006/0031481 A1, which discloses the monitoring of a plurality of service proxies by collecting and aggregating data and triggering events according to rules predefined by a user. Another example for a modeling tool for composite e-services and functionality, which allows monitoring the status of service execution, can be found in the US 7,222,334 B2.

[0007] In order to increase the reliability of the rules of a monitoring system, it is known from the prior art to use self-learning and rule-based monitoring processes as e.g. depicted in Fig. 1. To this end, a rule is first obtained in a learning phase, preferably in a learning loop, which results in a proposed rule for the SOA component to be monitored. The proposed rule may then be validated through a testing phase comprising any number of tests, which may be performed in a loop until the proposed rule passes all tests. Once all tests are passed, the proposed rule may be automatically marked as validated.

[0008] However, the above described learning/testing process may be inefficient and time-consuming, since the whole learning/testing process typically has to be undergone numerous times in the life-cycle of a SOA. For example when new components are to be added to the SOA landscape, they cannot be monitored right away, since the corresponding monitoring components each have to undergo the full learning/testing process in order to obtain the monitoring-rules. This makes the SOA either inflexible and inefficient (in case the new components are not used before validated rules have been "learned") or, more importantly, may impose a severe security-risk (in case the new components are already used although the rules are not yet validated).

[0009] In view of the above, it is therefore the technical problem of the present invention to improve the rule-based monitoring of a service-oriented architecture (SOA), thereby at least partly overcoming the disadvantages of the prior art.

3. Summary of the Invention

[0010] This problem is in one aspect solved by a method for the rule-based monitoring of a component added to a service-oriented architecture (SOA), the SOA comprising a plurality of components being monitored by a plurality of monitoring components. In the embodiment of claim 1, the method comprises the following steps:

a. selecting a component from the plurality of components of the SOA which is similar to the added component based on a similarity-metric;

b. selecting at least one rule from the monitoring component of the selected component, wherein the at least one rule has already been validated during the operation of the SOA; and

c. adding the at least one rule to a new monitoring component of the added component.

[0011] Accordingly, when adding a new component to a SOA, a similar existing component of the SOA is identified through a similarity-metric. If the identified component, i.e. its monitoring component, has rules which have already been validated in the operation of the SOA, these rules may be added to the monitoring component of the new component. The invention greatly simplifies the monitoring of a SOA when adding a new component by reusing already validated rules of similar monitoring components.

[0012] In one aspect of the invention, the similarity-metric may calculate a similarity-value indicating the similarity of the added component with the plurality of components of the SOA and the component with the highest similarity-value may be selected. Additionally or alternatively, the component may only be selected in step a. if the similarity-value is greater or equal to a pre-defined value. This ensures that only rules from sufficiently similar components are reused.

[0013] In another aspect, the similarity-metric may categorize the components into groups such as service, application server and / or computer hardware. Such a categorization serves to ensure that only components of similar type are compared by the similarity-metric.

[0014] If the components are services, the similarity-value may be calculated based on any of the group of information comprising the number and/or the names of methods provided by the services, their input parameter types and/or their output parameter types. Alternatively, if the components are application servers, the similarity-value may be calculated based on any of the group of information comprising their product types, product versions, used software and/or used hardware. Lastly, if the components are computer hardware, the similarity-value may be calculated based on any of the group of information comprising their operating systems, operating system versions and/or hardware configurations.

[0015] The existing monitoring component may have undergone a learning-phase to obtain the at least one rule. Furthermore, the selected monitoring component may have undergone a testing-phase to validate the at least one rule during the operation of the SOA. Accordingly, with the present method, rules can be obtained for new components to be monitored without having to go through the full learning-phase of the respective monitoring component if a similar component is already

present in the monitored landscape which has already undergone a quality approved learning-phase.

[0016] The method may further comprise the step of providing the new monitoring component and the at least one added rule to a testing-phase to validate the at least one rule during the operation of the SOA. This ensures that the newly added rules are tested and validated in the context of the added component. Additionally or alternatively, the method may further comprise the step of providing the new monitoring component and the at least one added rule to a learning-phase to determine at least one parameter of the at least one rule.

[0017] The present invention also relates to a monitoring system for the rule-based monitoring of a component added to a service-oriented architecture (SOA), the SOA comprising a plurality of components being monitored by a plurality of monitoring components, the monitoring system being adapted for performing any of the described methods. Furthermore, a computer program is provided comprising instructions for implementing any of the described methods.

4. Short Description of the Drawings

[0018] In the following detailed description presently preferred embodiments of the invention are further described with reference to the following figures:

Fig. 1: A schematic overview of a learning/testing process according to the prior art; and

Fig. 2: A schematic overview of an embodiment of the present invention;

4. Detailed Description

[0019] In the following, an exemplary embodiment of the monitoring method and system of the present invention will be described. The embodiment of the invention comprises a system and a method for a monitoring process which simplifies the obtaining of reliable rules for new components added to a SOA landscape by reusing existing validated rules of similar components of the SOA.

[0020] A SOA landscape typically comprises many components that can interact with each other in a complex way. Such components of a SOA can be of very different origin comprising objects such as services, especially web services, applications, clients, hardware and software components, and many more.

[0021] In order to monitor the components of a SOA, it is known to represent each of the components by a monitoring component which collects runtime data during the operation of the SOA. In this context, Fig. 2 presents an overview of an extremely simplified SOA landscape, which comprises a plurality of components C2, ..., Cn, which are accompanied by monitoring components M2, ..., Mn. In the depicted example, each of the monitoring components M2, ..., Mn monitors one correspond-

ing component C2, ..., Cn of the SOA. However, it is also possible to have one monitoring component for monitoring more than one of the components of the SOA. It is further possible to monitor only some of the components of the SOA.

[0022] The data collected by the monitoring components M2, ..., Mn may comprise actually measured values of the bandwidth, the message length of the messages issued by the components of the SOA, the number of messages per time, the duration or any other suitable value related to an operation of the components of the SOA. If the monitored component is e.g. a hardware component of the SOA, different values which are specific to this component may be monitored.

[0023] Combining the collected runtime data obtained by the monitoring components will provide an overall picture on how well the SOA landscape is operating. In order to filter the collected runtime data, i.e. to exclude irrelevant information, monitoring-rules may be attached to the monitoring components, so that the monitoring components only collect the relevant data from the components of the SOA landscape. As can be seen in Fig. 2, the monitoring component M2 has one or more monitoring-rules R2 attached. Furthermore, such monitoring-rules may also define what to do with the collected information, i.e. writing it to a log-file or issuing warning messages and / or alerts if certain boundaries are exceeded.

[0024] Monitoring-rules may be statically pre-defined e.g. by experienced administration personnel of the SOA. However, this is inflexible with respect to changes of the components of the SOA, since new rules would have to be formulated manually, which is time-consuming and error-prone. To this end, it is known to provide a self-learning monitoring process as depicted in Fig. 1. First, in a learning phase, a certain starting rule is incrementally refined, e.g. based on statistical values, until a proposed rule is obtained. This proposed rule is subsequently tested in a testing phase comprising any number of tests. When the proposed rule has passed all tests, it may be automatically or manually set to "validated" and used for the monitoring of the SOA. It should be appreciated that the learning of monitoring-rules could be accomplished in a test-run of the SOA as well as during its operational performance.

[0025] When a new component C1 (see Fig. 2) is to be added to the SOA, normally the whole learning/testing process would have to be undergone in order to obtain reliable rules for a monitoring component M1 corresponding to the new component C1, which is inefficient and time-consuming and might delay the operation of the SOA, because the new component C1 cannot be used productively until a reliable monitoring-rule has been "learned".

[0026] On the other hand, the new component C1 may in certain aspects be similar to or almost identical with the existing components C2, ..., Cn, e.g. two application servers of the same type running on two (nearly) identical computers or two web services that have nearly the same

operations. Consequently, rules of the existing monitoring components M2, ..., Mn might be re-used for the new monitoring component M1 in case the corresponding components are sufficiently similar.

[0027] This similarity may be automatically quantified by the present invention based on a similarity-metric (not shown in Fig. 2), which preferably returns a number between 0 and 100 indicating the amount of similarity between two components of the SOA. If the new component C1 to be added to the SOA has a certain required amount of similarity, e.g. 90% or more, with one of the existing components, e.g. with the component C2, it is likely that the validated rules R2 will also be similar to the rules needed to monitor the new component C1, if not identical.

[0028] Consequently, adding the already validated rules R2 to a monitoring component M1 of the new component C1 (illustrated by the horizontal arrow in Fig. 2) will save time and efforts and considerably speed up the start of the monitoring process for a changed or extended SOA landscape, since the learning/testing process does not have to be undergone.

[0029] The similarity-value is preferably calculated pair-wise for the new component C1 with each of the existing components C2, ..., Cn and the existing component with the highest similarity value is selected. In order to restrict the comparison of two components to reasonable combinations, only components from the same group may be compared, i.e. only web services or only hardware components of the same type, etc.. The similarity-value may be calculated in a number of ways depending on the types of the compared components, for example:

- If the components are services (e.g. web services), the similarity-value may be calculated by comparing the number, i.e. the amount, and the names of the operations, i.e. sub-programs, provided by the services and/or by comparing their input and output parameter types, such as integer, string or floating-point number.
- If the components are application servers, the similarity-value may be calculated by comparing their product types, product versions and/or their used software (e.g. operating system, runtime environment) or hardware (e.g. available memory).
- If the components are computer hardware, the similarity-value may be calculated by comparing their operating systems, operating system versions and/or hardware configurations (e.g. main memory, hard disk space, types of hardware components such as graphic cards or motherboards).

[0030] Furthermore, in that preferably only operationally validated rules are re-used for new monitoring components, i.e. rules that already have undergone the full testing-phase, a sufficient quality and robustness of the reused rules is guaranteed.

[0031] To further ensure the feasibility of the re-used

rules, the new monitoring component M1 may enter the testing-phase in order to test the re-used rule R2 for the monitoring of the component C1. This approach has a considerably improved efficiency compared to the prior art approach, since the costly learning-phase of the learning/testing process can be skipped. If the similarity is sufficiently high, e.g. if two identical computer hardware components are selected, even the testing-phase may be skipped in order to increase the efficiency even more. On the other hand, the learning-phase may in some scenarios still be undergone by the new monitoring component M1, e.g. if the added rules R2 comprise parameters and the concrete parameter values have to be determined for the new monitoring component M1.

Claims

1. A method for the rule-based monitoring of a component (C1) added to a service-oriented architecture (SOA), the SOA comprising a plurality of components (C2, ..., Cn) being monitored by a plurality of monitoring components (M2, ..., Mn), the method comprising the following steps:
 - a. selecting a component (C2) from the plurality of components (C2, ..., Cn) of the SOA which is similar to the added component (C1) based on a similarity-metric;
 - b. selecting at least one rule (R2) from the monitoring component (M2) of the selected component (C2), wherein the at least one rule (R2) has already been validated during the operation of the SOA; and
 - c. adding the at least one rule (R2) to a new monitoring component (M1) of the added component (C1).
2. The method of claim 1, wherein the similarity-metric calculates a similarity-value indicating the similarity of the added component (C1) with the plurality of components (C2, ..., Cn) of the SOA and wherein the component (C2) with the highest similarity-value is selected.
3. The method of claim 2, wherein the component (C2) is only selected in step a. if the similarity-value is greater or equal to a pre-defined value.
4. The method of any of the preceding claims, wherein the similarity-metric categorizes the components (C1, C2,..., Cn) into groups such as service, application server and / or computer hardware.
5. The method of any of the preceding claims 2 - 4, wherein the components (C1, C2,..., Cn) are services and wherein the similarity-value is calculated based on any of the group of information comprising
 - the number and/or the names of operations provided by the services, their input parameter types and/or their output parameter types.
6. The method of any of the preceding claims 2 - 4, wherein the components (C1, C2,..., Cn) are application servers and wherein the similarity-value is calculated based on any of the group of information comprising their product types, product versions, used software, and/or used hardware.
7. The method of any of the preceding claims 2 - 4, wherein the components (C1, C2,..., Cn) are computer hardware and wherein the similarity-value is calculated based on any of the group of information comprising their operating systems, operating system versions and/or hardware configurations.
8. The method of any of the preceding claims, wherein the selected monitoring component (M2) has undergone a learning-phase to obtain the at least one rule (R2).
9. The method of any of the preceding claims, wherein the selected monitoring component (M2) has undergone a testing-phase to validate the at least one rule (R2) during the operation of the SOA.
10. The method of any of the preceding claims, further comprising the step of providing the new monitoring component (M1) and the at least one added rule (R2) to a testing-phase to validate the at least one rule (R2) during the operation of the SOA.
11. The method of any of the preceding claims, further comprising the step of providing the new monitoring component (M1) and the at least one added rule (R2) to a learning-phase to determine at least one parameter of the at least one rule (R2).
12. A monitoring system for the rule-based monitoring of a component (C1) added to a service-oriented architecture (SOA), the SOA comprising a plurality of components (C2, ..., Cn) being monitored by a plurality of monitoring components (M2, ..., Mn), the monitoring system being adapted for performing a method of any of the claims 1 - 11.
13. A computer program comprising instructions for implementing a method of any of the claims 1 - 11.

Fig. 1 (Prior Art)

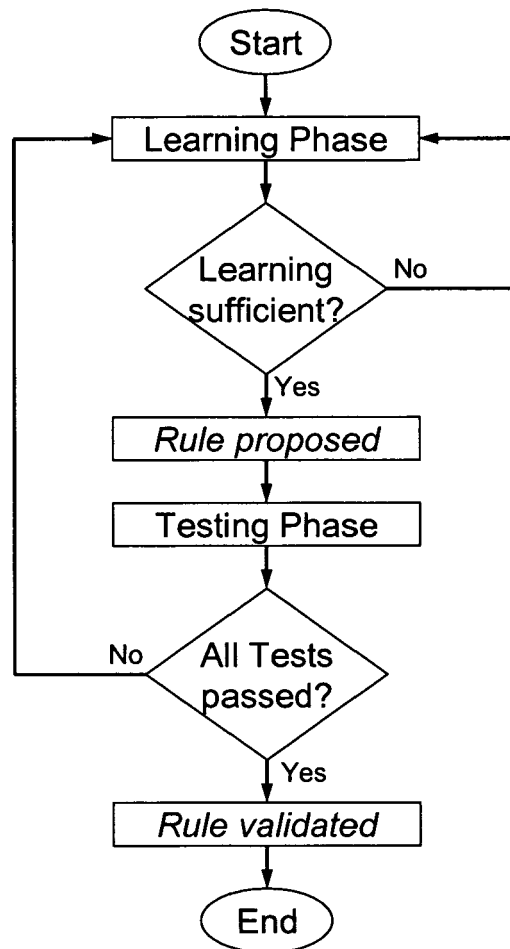



Fig. 2





EUROPEAN SEARCH REPORT

Application Number
EP 08 01 7502

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
Y	GEHLEN G ET AL: "A Rule based Data Monitoring Middleware for Mobile Applications" VEHICULAR TECHNOLOGY CONFERENCE, 2005. VTC 2005-SPRING. 2005 IEEE 61ST, IEEE, PISCATAWAY, NJ, USA, vol. 5, 30 May 2005 (2005-05-30), pages 2910-2914, XP010855982 ISBN: 978-0-7803-8887-1 * paragraph [III.A] *	1-13	INV. G06N5/02
Y	US 5 317 677 A (DOLAN CHARLES P [US] ET AL) 31 May 1994 (1994-05-31) * abstract *	1-13	
D,A	US 2006/031481 A1 (PATRICK PAUL B [US] ET AL) 9 February 2006 (2006-02-09) * abstract *	1-13	
D,A	EP 1 863 258 A (SOFTWARE AG [DE]) 5 December 2007 (2007-12-05) * abstract *	1-13	TECHNICAL FIELDS SEARCHED (IPC) G06N G06F
A	ROMAN VACULIN ET AL: "Semantic Web Services Monitoring: An OWL-S Based Approach" HAWAII INTERNATIONAL CONFERENCE ON SYSTEM SCIENCES, PROCEEDINGS OF THE 41ST ANNUAL, IEEE, PISCATAWAY, NJ, USA, 1 January 2008 (2008-01-01), page 313, XP031208183 ISBN: 978-0-7695-3075-8 * paragraph [0004] *	1-13	
The present search report has been drawn up for all claims			
Place of search The Hague		Date of completion of the search 20 May 2009	Examiner Valencia, Erika
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

6
EPO FORM 1503 03.82 (P04C01)



EUROPEAN SEARCH REPORT

Application Number
EP 08 01 7502

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
A	<p>XIANGJI HUANG ET AL: "Discovery of interesting association rules from Livelink web log data"</p> <p>DATA MINING, 2002. PROCEEDINGS. 2002 IEEE INTERNATIONAL CONFERENCE ON MAEBASHI CITY, JAPAN 9-12 DEC. 2002, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 9 December 2002 (2002-12-09), pages 763-766, XP010804978</p> <p>ISBN: 978-0-7695-1754-4</p> <p>* paragraph [0004] *</p> <p>-----</p>	1-13	
			TECHNICAL FIELDS SEARCHED (IPC)
The present search report has been drawn up for all claims			
Place of search The Hague		Date of completion of the search 20 May 2009	Examiner Valencia, Erika
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

6

EPO FORM 1503 03.82 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 08 01 7502

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

20-05-2009

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5317677	A	31-05-1994	NONE	

US 2006031481	A1	09-02-2006	NONE	

EP 1863258	A	05-12-2007	AT 421221 T	15-01-2009
			CN 101087312 A	12-12-2007
			US 2007300240 A1	27-12-2007

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- EP 1863258 A [0005]
- US 20060031481 A1 [0006]
- US 7222334 B2 [0006]